The Tenth International Symposium on Operations Research and Its Applications (ISORA 2011)
Dunhuang, China, August 28–31, 2011

An Approximation Algorithm For The Stochastic Fault-tolerant Facility Location Problem

Chen-Chen Wu[1] Da-Chuan Xu[1,2] Jia Shu[3]

[1]Department of Applied Mathematics, Beijing University of Technology
Beijing 100124, P.R. China. Email: wuchenchen@emails.bjut.edu.cn; xudc@bjut.edu.cn.
[2]Corresponding author.
[3]Department of Management Science and Engineering, School of Economics and Management
Southeast University, Nanjing 210096, P.R. China. Email: tliapshujia@yahoo.com.

Abstract In this paper, we study a stochastic version of the fault-tolerant facility location problem. By exploiting the stochastic structure, we propose a 5-approximation algorithm which uses the LP-rounding technique based on the revised optimal solution to the linear programming relaxation of the stochastic fault-tolerant facility location problem.

Keywords Facility location problem; Approximation algorithm; LP rounding

1 Introduction

The facility location problem (FLP) has been widely studied in the operations research literature. The first constant approximation ratio is 3.16 which is given by Shomys et al. [15] using the LP-rounding technique. Then, there are several results with respect to the approximation algorithm (cf. [7, 8, 11]) for this problem. The current best approximation ratio is 1.488 by Li [16] using LP-rounding. Sviridenko [18] shows that it is impossible to design an approximation algorithm with an approximation ratio smaller than or equal to 1.463, unless P=NP. For a discussion of the variants of the FLP, we refer to [5, 21] and the references therein.

The fault-tolerant facility location problem (FTFLP) is one of the most important variants of the FLP. In the setting of the FTFLP, each client is required to be assigned to more than one facility in order to prevent that the open facility maybe failure which leads to the situation that some clients cannot be served. The assignment cost of each client is a weighted combination of the distances to the facilities to which the client connects. This problem is first studied by Jain and Vazirani [9], in which they give a primal-dual algorithm with a logarithmic of the largest requirement approximation ratio for the weight of each client assigning to distinct facilities being uniform. Then, Swamy and Shmoys [14] give a 2.076-approximation algorithm for the uniform weight. The approximation ratio is further improved to 1.7245 by Byrka et al. [3]. Guha et al. [6] give a 2.408-approximation algorithm for the generalized non-uniform weight case. On the other hand, the stochastic facility location problem (SFLP) has also been studied extensively in the literature.

Generally speaking, the SFLP contains a 2-stage stochastic process: in the first stage of the process, the information of the clients are unknown. The possible scenarios and their corresponding probability distributions are given in the second stage. Each facility at different stages and scenarios is different. We can open some facilities in the first stage which can serve any client in any scenario and the facilities open in each scenario can only be used to serve clients in that scenario. The problem is introduced by Ravi and Sinha [12], who present an LP-rounding 8-approximation algorithm. Then, there are several results related to the SFLP (cf. [10, 13, 17]). The current best approximation ratio for the SFLP is 1.8526 by Ye and Zhang [19].

In this paper, we are interested in the stochastic fault-tolerant facility location problem (SFTFLP) in which each client in each scenario is specified to be assigned to more than one facility. Intuitively, some facilities could be failure so that the clients require some backups. We give a LP-rounding 5-approximation algorithm by integrating the techniques of [6, 12, 15].

2 The 2-stage stochastic fault-tolerant facility location problem

For the 2-stage stochastic fault-tolerant facility location problem, we are given a facility set F in the first stage only. For each facility i, its open cost is f_i^0 in the first stage. In the second stage, we are given the client sets D_s which need to be served for each scenario $s \in \{1,2\cdots,S\}$, a probability p_s associated with each scenario s, and a distance function $c : F \times (\bigcup_s D_s) \longrightarrow R_+$ which is metric, i.e., satisfies symmetry, nonnegativity, and the triangle inequality. In each scenario s, it is allowable to open some facilities to serve the clients in D_s. The open cost of facility i in scenario s is f_i^s. In each scenario s, each client in D_s need to be assigned to r_j distinct facilities (according to certain weights) which are opened only in the first stage and the corresponding scenario. Let the weights of assigning j to the r_j facilities be $w_j^1 \geq w_j^2 \geq ... \geq w_j^{r_j}$, that is, the assignment cost of client j is the weighted combination of the distance to r_j closest facilities. The goal of the problem is to assign the clients to the opened facilities such that the total expected facility open and assignment costs are minimized.

We denote $\mathscr{F} := \{(i,t)|i \in F, t = 0,1,...,S\}$, $\mathscr{D} := \{(j,r,s)|s = 1,...,S, j \in D_s, r = 1,...,r_j\}$, and $\tilde{\mathscr{D}} := \{(j,s)|s = 1,...,S, j \in D_s\}$. Also, let $p_0 := 1$. We call the element in $\mathscr{F}$ a facility-scenario pair and similarly, the elements in $\mathscr{D}$ and $\tilde{\mathscr{D}}$ are called a client-copy-scenario triple and a client-scenario pair, respectively. For each facility-scenario pair (j,s), $r \in 1,2,\cdots,r_j$ is called a copy of (j,s). Following the above notations, we can restate the problem as: given a facility-scenario pair set $\mathscr{F}$ and a client-copy-scenario triple set $\mathscr{D}$, we intend to open a facility-scenario pair set F_0 in the first stage and F_s in scenario s of stage 2 in which the facility open cost is $p_s f_i^s$, and assign each client-scenario pair to r_j distinct opened facility-scenario pair in which each client-copy-scenario triple (j,r,s) can only be assigned to the facility-scenario pairs in F_0 and F_s, so that the total cost including the facility open cost and the assignment cost is minimized. In order to ensure (j,r,s) cannot be assigned to the facility-scenario pair (i,t), where $t \neq 0$ or s, we define the distance between a facility-scenario pair and a client-scenario pair as follows:

$$c_{ij}^{ts} = \begin{cases} c_{ij} & \text{if } t = 0 \text{ or } s; \\ +\infty & \text{otherwise.} \end{cases}$$

To this end, the problem can be formulated as the following integer program:

$$
(IP) \qquad
\begin{array}{ll}
\min & \sum\limits_{(i,t)\in\mathscr{F}} p_t f_i^t y_i^t + \sum\limits_{(i,t)\in\mathscr{F},(j,r,s)\in\mathscr{D}} p_s w_j^r c_{ij}^{ts} x_{ij}^{tsr} \\
\text{s.t.} & \sum\limits_{(i,t)\in\mathscr{F}} x_{ij}^{tsr} \geq 1, \qquad \forall (j,r,s)\in\mathscr{D}, \\
& \sum\limits_{r=1}^{r_j} x_{ij}^{tsr} \leq y_i^t, \qquad \forall (i,t)\in\mathscr{F}, \forall (j,r,s)\in\mathscr{D}, \\
& x_{ij}^{tsr} \in \{0,1\}, \qquad \forall (j,r,s)\in\mathscr{D}, \\
& y_i^t \in \{0,1\}, \qquad \forall (i,t)\in\mathscr{F},
\end{array}
\qquad (1)
$$

in which, y_i^t indicates whether the facility-scenario pair (i,t) (including $t=0$) is open; x_{ij}^{tsr} denotes whether client j in scenario s is assigned to the facility-scenario pair (i,t) and (i,t) is the rth closest open facility-scenario pair to j. The first constraint of (1) requires each client-copy-scenario pair (j,r,s) should be assigned to a facility-scenario pair. The second constraint of (1) models that if the client-copy-scenario triple (j,r,s) is assigned to the facility-scenario pair (i,t), the pair (i,t) should open, and each facility-scenario pair (i,t) can only serve one copy of the same client-scenario pair (j,s).

By relaxing the integrality constraints, we obtain the following LP relaxation:

$$
(LP) \qquad
\begin{array}{ll}
\min & \sum\limits_{(i,t)\in\mathscr{F}} p_t f_i^t y_i^t + \sum\limits_{(i,t)\in\mathscr{F},(j,r,s)\in\mathscr{D}} p_s c_{ij}^{ts} w_j^r x_{ij}^{tsr} \\
\text{s.t.} & \sum\limits_{(i,t)\in\mathscr{F}} x_{ij}^{tsr} \geq 1, \qquad \forall (j,r,s)\in\mathscr{D} \\
& \sum\limits_{r=1}^{r_j} x_{ij}^{tsr} \leq y_i^t, \qquad \forall (i,t)\in\mathscr{F}, \forall (j,r,s)\in\mathscr{D} \\
& x_{ij}^{tsr} \geq 0, \qquad \forall (j,r,s)\in\mathscr{D} \\
& 0 \leq y_i^t \leq 1, \qquad \forall (i,t)\in\mathscr{F}.
\end{array}
\qquad (2)
$$

Let $F^* = \sum\limits_{(i,t)\in\mathscr{F}} p_t f_i^t y_i^t$ and $C^* = \sum\limits_{(i,t)\in\mathscr{F}} \sum\limits_{(j,r,s)\in\mathscr{D}} p_t w_j^r c_{ij}^{ts} x_{ij}^{ts}$ be the optimal fractional facility cost and assignment cost, respectively.

3 The algorithm

Now we proceed to describe the algorithm as follows.

Algorithm 1. *(***LP-rounding algorithm***)*

Step 1. Solving the LP relaxation and constructing a consistent solution.

Solve the LP relaxation (2) to obtain the optimal solution (x,y). We next convert it to another solution $(\bar{x},y)$ which has some useful properties as follows. For each client-scenario pair $(j,s)\in\mathscr{D}$, sort all facility-scenario pairs according to their distances to (j,s) in a nondecreasing order. For the same facility in the first stage and any scenario, we put the first stage facility before the scenario one. Then, we assign the first copy of (j,s) to the facility-scenario pair (i,t)($t=0$ or s) in terms of the above ordering, i.e., we set $\bar{x}_{ij}^{ts1} := y_i^t$. We repeat the above operations until the summation of $\bar{x}$ equals to 1. Therefore, the last facility-scenario pair (i,t) may not be used completely, i.e., $\bar{x}_{ij}^{ts1} < y_i^t$. For the second copy, we set $\bar{x}_{ij}^{ts2} := y_i^t - \bar{x}_{ij}^{ts1}$. After picking up one unit of the facility-scenario pair, we turn to the next copy. Repeat *this process for all copies of the pairs. Then, we can obtain a solution $(\bar{x}, y)$ which is called a* **consistent solution***.*

Step 2. Filtering and scaling.

For each client-scenario pair (j,s), we run the following operations in the increasing order of copies $r = 1,2,\cdots,r_j$. For each client-copy-scenario triple (j,r,s), sort the facility-scenario pairs in the nondecreasing order of their distances to (j,r,s)(same as Step 1). Then, we set $C_j^{sr} := c_{i^*j}^{t^*sr}$, where (i^*,t^*) is the first pair such that $\sum_{(i,t):c_{ij}^{ts} \le c_{i^*j}^{t^*s}} \bar{x}_{ij}^{tsr} \ge \frac{2}{5}$. We also set

$$\hat{x}_{ij}^{tsr} := \begin{cases} \bar{x}_{ij}^{tsr} & \text{if } c_{ij}^{ts} < C_j^{sr}; \\ \frac{2}{5} - \sum_{(i,t):c_{ij}^{ts} < c_{i*j}^{t*s}} \hat{x}_{ij}^{tsr} & \text{if } c_{ij}^{ts} = C_j^{sr}; \\ 0 & \text{otherwise.} \end{cases}$$

Furthermore, we scale $\hat{x}_{ij}^{tsr}$ by $\frac{5}{2}$ to obtain $\tilde{x}_{ij}^{tsr}$ such that $\sum_{(i,t)} \tilde{x}_{ij}^{tsr} = 1$. Therefore, we let $\tilde{y}_i^t := \min\{\frac{5}{2}y_i^t, 1\}$. Let $\bar{C}(j,r,s)$ denote the facility-scenario pair set which serves (j,r,s) according to the solution $(\tilde{x}, \tilde{y})$.

Step 3. Clustering.

For ease of exposition, we denote $F_0(j,r,s) := \{(i,0)|\tilde{x}_{ij}^{0sr} > 0\}$, $F_s(j,r,s) := \{(i,s)| \tilde{x}_{ij}^{ssr} > 0\}$, $w_0(j,r,s) := \sum_{(i,0)\in F_0(j,r,s)} \tilde{x}_{ij}^{0sr}$, and $w_s(j,r,s) := \sum_{(i,s)\in F_s(j,r,s)} \tilde{x}_{ij}^{ssr}$. Let $\bar{\mathscr{D}}$ denote the set of the unassigned client-copy-scenario triples which will be updated in the clustering process. Initially, $\bar{\mathscr{D}} := \mathscr{D}$ which contains all the client-copy-scenario triples. For each client-copy-scenario triple (j,r,s), let $\bar{C}(j,r,s)$ denote the facility-scenario pair set which serves (j,r,s) according to the solution $(\tilde{x}, \tilde{y})$.

Step 3.1 Picking center.

Arrange all the client-copy-scenario triples in the nondecreasing order of the value of C_j^{sr}. Assume (j,r,s) has the smallest C_j^{sr}, which we call a center.

Step 3.2 Choosing facility-scenario pair set.

We pick up the pairs in $F_0(j,r,s)$, if $w_0(j,r,s) \ge w_s(j,r,s)$; otherwise pick up the pairs in $F_s(j,r,s)$. Assume the chosen facility-scenario pair is (i,t)(t is 0 or s). Let $N_F(j,r,s)$ and Y be the facility-scenario pair set which will be assigned to the center (j,r,s) and the sum of the values of the facility-scenario pairs in $N_F(j,r,s)$. Given the value of Y, **partition** *(i,t) means partitioning (i,t) into (i_1,t) and (i_2,t) such that $\tilde{y}_{i_1}^t := w_t(j,r,s)$ and $\tilde{y}_{i_2}^t := Y - \tilde{y}_{i_2}^t$, while maintaining $\sum_r \tilde{x}_{ij}^{tsr} \le \tilde{y}_i^t$ for both $i = i_1$ and $i = i_2$.*

Initially, set $N_F(j,r,s) := \emptyset$ and $Y := 0$. Let $(i',t) := \arg\min_{(i,t)\in F_t(j,r,s)} p_t f_i^t$ and set $Y := \tilde{y}_{i'}^t$, $F_t(j,r,s) := F_t(j,r,s) - \{(i',t)\}$, and $N_F(j,r,s) := \{(i_1,t)\}$. If $Y > w_t(j,r,s)$, **partition** *(i',t). Otherwise, we check whether Y is exactly $w_t(j,r,s)$ or not. If $Y = w_t(j,r,s)$, we can obtain a facility-scenario pair $N_F(j,r,s)$. If $Y < w_t(j,r,s)$, we choose a facility-scenario pair (i,t) in $F_t(j,r,s)$ arbitrarily and set $Y := Y + \tilde{y}_i^t, F_t(j,r,s) := F_t(j,r,s) - \{(i,t)\}$. After that, it is possible that $Y > w_t(j,r,s)$. For this,* **partition** *(i,t) and set $N_F(j,r,s) := N_F(j,r,s) \cup \{(i_1,t)\}$, $Y := Y - \tilde{y}_{i_2}^t$. Otherwise, set $N_F(j,r,s) := N_F(j,r,s) \cup \{(i,t)\}$ and check whether the value of Y is bigger than $w_t(j,r,s)$ again as above.*

At the end of this step, we can obtain a facility-scenario pair set $N_F(j,r,s)$ of a center (j,r,s) whose value is $w_t(j,r,s)$ exactly.

Step 3.3 Choosing client-copy-scenario triple set and reassigning.

For each remaining client-scenario pair (j',s'), we denote

$$R(j',s') := \{r|\exists(i,t) \in N_F(j,r,s), \text{ s.t. } \tilde{x}_{ij'}^{ts'r} > 0\}.$$

Assume $R(j',s') = \{r_1, r_2, \cdots, r_{k(j',s')}\}$. Then, $N_D(j,r,s) := \bigcup_{R(j',s')\neq\emptyset} \{(j',r_1,s')\}$ and $\bar{\mathscr{D}} := \bar{\mathscr{D}} - N_D(j,r,s)$.

It is necessary to deal with the remaining copies of (j',s'). Initially, $T(j',r_1,s')$ is the facility-scenario pair set which serves (j',r_1,s') but not in $N_F(j,r,s)$, that is, $T(j',r_1,s') := \bar{C}(j',r_1,s') - N_F(j,r,s)$. For a client-copy-scenario triple (j',r_m,s'), set $X := \sum_{(i,t)\in N_F(j,r,s)} \tilde{x}_{ij'}^{ts'r_m}$. Then, check whether X is exactly 0 or not. If it is true, we turn to the next copy. Otherwise, let $(\tilde{i},\tilde{t}) := \arg\min_{(i,t)\in T(j',r_1,s')} c_{ij'}^{ts'r_1}$. If $X > \tilde{y}_{\tilde{i}}^{\tilde{t}}$, set $\tilde{x}_{\tilde{i}j'}^{\tilde{t}s'r_m} := \tilde{y}_{\tilde{i}}^{\tilde{t}}$, $T(j',r_1,s') := T(j',r_1,s') - \{(\tilde{i},\tilde{t})\}$, and $X := X - \tilde{x}_{\tilde{i}j'}^{\tilde{t}s'r_m}$; Otherwise, set $\tilde{x}_{\tilde{i}j'}^{\tilde{t}s'r_m} := X$, $\tilde{y}_{\tilde{i}}^{\tilde{t}} := \tilde{y}_{\tilde{i}}^{\tilde{t}} - X$, and $X := X - \tilde{x}_{\tilde{i}j'}^{\tilde{t}s'r_m}$. Finally, for all $(i,t) \in N_F(j',r_1,s')$, set $\tilde{x}_{ij'}^{ts'r_m} := 0$.

Step 3.4 Constructing cluster.

We call $N_F(j,r,s)$ and $N_D(j,r,s)$ combined with (j,r,s) is a cluster centered at (j,r,s).

Repeat the above procedure until each client-copy-scenario triple belongs to some cluster.

Step 4. Rounding.

We can obtain several disjoint clusters at the end of Step 3. In each cluster, we open the cheapest facility-client pair and close the others, and assign each client-copy-scenario triple to the opened facility-scenario pair.

Theorem 1. *Algorithm 1 is a well-defined polynomial-time 5-approximation algorithm for the SFTFLP.*

A thorough proof of Theorem 1 will appear in the full version of the paper.

Acknowledgements

The second author's research is supported by Natural Natural Science Foundation of China (No. 11071268, 60773185), Beijing Natural Science Foundation (No. 1102001), and PHR(IHLB). The third author's research is supported by National Natural Science Foundation of China (No. 70801014) and Program for New Century Excellent Talents in University of China (No. NCET-09-0292).

References

[1] A. Ageev, Y. Ye and J. Zhang, Improved combinatorial approximation algorithms for the k-level facility location problem, *SIAM Journal on Discrete Mathematics*,18(2004)207-217.

[2] J. Byrka and K.I. Aardal An optimal bifactor approximation algorithm for the metric uncapacitated facility location problem, *SIAM Journal on Computing*, 39(2010)2212-2231.

[3] J. Bykra, A. Srinivasam and C. Swamy, Fault-tolerant Facility Location: a randomized dependent LP-rounding algorithm, *Proceedings of the 14th Conference on Integer Programming and Combinatorial Optimization*, 2010, pp. 244-257.

[4] X. Chen and B. Chen, Approximation algorithms for soft-capacitated facility location in capacitated network design, *Algorithmica*, 53(2007)263-297.

[5] S. Guha and S. Khuller, Greedy strike back: Improved facility location algorithms, *Journal of Algorithms*, 31(1999)228-248.

[6] S. Guha, A. Meyerson and K. Munagala, A constant factor approximation algorithm for the fault-tolerant facility location problem, *Journal of Algorithms*, 48(2003)429-440.

[7] K. Jain, M. Mahdian, E. Markakis, A. Saberi and V.V. Vazirani, Greedy facility location algorithm analyzed using dual fitting with factor-revealing LP, *Journal of the ACM*, 50(2003)795-824.

[8] K. Jain and V.V. Vazirani, Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and Lagrangian relaxation, *Journal of the ACM*, 48(2001)274-296.

[9] K. Jain and V.V. Vazirani, Approximation algorithms for fault tolerant metric facility location problem, *Algorithmica*, 38(2003)433-439.

[10] M. Mahdian, Facility location and the analysis of algorithms through factor-revealing programs, Ph. D. thesis, MIT, Cambridge, MA, 2004.

[11] M. Mahdian, Y. Ye and J. Zhang, Approximation algorithms for metric facility location problems, *SIAM Journal on Computing*, 36(2006)411-432.

[12] R. Ravi and A. Sinha, Hedging uncertainty: approximation algorithms for stochastic optimization problems, *Mathematical Programming*, 108(2006)97-114.

[13] D.B. Shmoys and C. Swamy, An approximation scheme for stochastic linear programming and its application to stochastic integer programs, *Journal of the ACM*, 53(2006)978-1012.

[14] D.B. Shmoys and C. Swamy, Fault-tolerant facility location, *ACM Transactions on Algorithms*, 4(2008)article No.51.

[15] D.B. Shmoys, E. Tardös and K.I. Aardal, Approximation algorithms for facility location problems, *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, 1997, pp. 265-274.

[16] S. Li, A 1.488-approximation algorithm for the uncapacitated facility location problem, to appear in: *Proceedings of the 38th International Colloquium on Automata, Languages and Programming*, 2011.

[17] A. Srinivasan, Approximation algorithms for stochastic and risk-averse optimization, *Proceedings of the 18th annual ACM-SIAM Symposium on Discrete Algorithms*, 2007, pp. 1305-1313.

[18] M. Sviridenko, An Improved approximation algorithm for the metric uncapacitated facility location problem, *Proceedings of the 9th Conference on Integer Programming and Combinatorial Optimization*, 2002, pp. 240-257.

[19] Y. Ye and J. Zhang, An approximation algorithm for the dynamic facility location problem, In Cheng M, Li Y, Du D(Eds), *Combinatorial Optimization in Communication Networks*, Kluwer Academic Publishers, 2005, pp. 623-637.

[20] J. Zhang, Approximating the two-level facility location problem via a quasi-greedy approach, *Mathematical Programming*, 108(2006)159-176.

[21] P. Zhang, A new approximation algorithm for the k-facility location problem, *Theoretical Computer Science*, 384(2007)126-135.